
Worldnet

Peter J. Denning

5 Jul 89

RIACS Technical Report TR-89.26

NASA Cooperative Agreement Number NCC 2-387

(NASA-CR-188845) WORLDNET (Research Inst.
for Advanced Computer Science) 15 pCSCL 05B

N92-10711

G3/82 Unclass
0043031

RIACS

Research Institute for Advanced Computer Science
An Institute of the Universities Space Research Association

IN-82

DATE OVERRIDE

43031

P-15

Worldnet

Peter J. Denning

Research Institute for Advanced Computer Science
NASA Ames Research Center

RIACS Technical Report TR-89.26
5 Jul 89

The expanding use of powerful workstations coupled to ubiquitous networks is transforming scientific and engineering research and the ways organizations around the world do business. By the year 2000, few enterprises will be able to succeed without mastery of this technology, which will be embodied in an information infrastructure based on a worldwide network. A recurring theme in all the discussions of what might be possible within the emerging Worldnet is people and machines working together in new ways across distance and time. This essay reviews the basic concepts on which the architecture of Worldnet must be built: coordination of action, authentication, privacy, and naming. Worldnet must provide additional functions to support the ongoing processes of suppliers and consumers: help services, aids for designing and producing subsystems, spinning off new machines, and resistance to attack. This discussion begins to reveal the constituent elements of a theory for Worldnet, a theory focused on what-people-do-with-computers rather than on what-computers-do.

This is a preprint of the column *The Science of Computing* for
American Scientist 77, No. 5 (September-October 1989).

Work reported herein was supported in part by Cooperative Agreement NCC 2-387
between the National Aeronautics and Space Administration (NASA)
and the Universities Space Research Association (USRA).

~~CONFIDENTIAL~~ INTENTIONAL RELEASE

Worldnet

Peter J. Denning

Research Institute for Advanced Computer Science

5 Jul 89

Business and government are moving inexorably toward electronic interdependence. Organizations are forming relationships across international boundaries that were not possible five years ago. Cellular telephone and fax are expanding worldwide. Science and engineering research now depends on powerful workstations attached to high-speed networks, encouraging collaboration and permitting access to remote resources (1). From all this is emerging a worldwide network of computers, which I will call Worldnet.

The components of Worldnet are computers, workstations, networks, and software, a mixture sometimes called information technology (2). By the year 2000, Worldnet will be ubiquitous and pervasive. It will be as important for conducting business, distributing information, and coordinating work as are the existing transportation and telecommunication networks. Few enterprises, commercial or scientific, will succeed without mastery of this technology.

Dramatic shifts in business and science will be produced by Worldnet. For example, most markets will be global; corporations will routinely conduct international business by network. The recorded history of a project, an organization, or a discipline will be available on-line, and new entries in those histories will be automatically created as people take action. Mass production of identical items will give way to production of items tailored to individual preferences, and network "boutiques" that specialize in customized products will be common. Within an organization, researchers, engineers, salespeople, and manufacturing experts will cooperate with users on new designs, making it possible to bring new products to market within two or three years of their conception. Subsystems to provide new objects and services will be routinely spun off as autonomous agents in Worldnet.

Large scientific and engineering projects will work effectively over long durations and large distances. New collaborations will arise because distance will no longer be a factor either in carrying out tasks or in sharing data: advanced scientific workstations will have audio and video monitors and screen-sharing protocols to support collaborations with distant colleagues. Individuals and organizations will have access to machines that can supply information in selected domains and help them locate and use resources. Data streams produced by instruments and sensors around the world will be brought together and the new findings distributed within a community. Thus it will be possible to achieve effective coordination of worldwide efforts such as averting famines, fighting AIDS, mapping the human genome, or modeling global climatic changes.

These developments will have profound effects on individuals as well. Computers with cellular telephone and fax connections will be common, enabling people to maintain

a link into the Worldnet no matter where they are -- at work, at home, or traveling.

Business trips and commuting will be much less of an interruption than today. Shopping from world markets will be common. Many people will work at home, linked fully with their associates through Worldnet.

The "information infrastructure" that makes all this possible will include networks and connections among them, protocols, and standards for network use; it will provide hookups, accounting, billing, maintenance, repair, and reconfiguration; it will supply directories of accessible users and resources; and it will provide a means by which a large variety of organizations can offer support services for users -- for example, news services, brokering, network advertising, and access to databases. A bill has been introduced in the Senate to stimulate the construction of this infrastructure in the research community with the aim of using it to support US leadership in high-performance computing, manufacture of high-performance computers, and applications in key disciplines. Although the focus of national policy will be on science and technology, the benefits to business and commerce are clear and immediate. The report of the MIT Commission on Industrial Productivity further underscores the importance of information technology to the future of US productivity (3,4); a similar conclusion has been reached by many third-world entrepreneurs who are working to bring their countries into the world business community as full partners within the next generation.

Against this background, I will speculate in the rest of this essay on what functions must be present in Worldnet. My speculations are grounded in an analysis of the fundamental actions that arise in all of the domains discussed above. They are designed to examine what people might do with computers, rather than what computers might do,

and thus they point toward a theoretical basis for Worldnet.

Underlying all the trends noted above is a recurrent theme of people and machines working together in new ways and across distance and time. Therefore, the most fundamental characteristic of the architecture of Worldnet must be support for coordination of action. What actions are to be coordinated? Among whom? The range of actions includes anything people can do in conversations together -- make requests and promises, speculate, work on shared documents, build new systems -- depending on the domain in which they are working. It also includes actions by machines that perform predetermined tasks. Because people and machines are both capable of initiating and carrying out actions, I use the term "agents" either for people or for machines in Worldnet.

Terry Winograd and Fernando Flores have formulated a model, called "conversation for action," of the underlying process by which humans use language to coordinate their actions (5,6). Anatol Holt has formulated a model of the processes of coordination within business organizations (7,8). Both models are already embodied in commercial software packages, and they could readily be part of the design of Worldnet.

In everyday business, we must frequently assess whether agents who have made promises to us are sincere and are competent to fulfill their promises. We avoid doing business with agents we do not trust. The processes by which trust is established include repeated direct observation of satisfactory performance and certifications by competent authorities (or machines) that we already trust. The processes by which we assure ourselves that an agent is one previously identified as trustworthy, collectively called authentication, include recognition of familiar faces, voices, or signatures, login

protocols on computers, and cryptographic protocols. Audio-video links between advanced workstations will help support authentication. In networks, where messages can be replayed, we may need to reestablish authentication repeatedly throughout a conversation.

Four examples illustrate the intimate connection between authentication and effective action. First, after you make private information available to a trusted agent, you want assurances that the agent will not grant access to others. Second, an agent holding a document or certificate that confers a particular authority needs to be able to prove the document's authenticity in case of a challenge. Third, when opening a network connection to a remote agent, you need verification that the agent is actually the one named. Fourth, when working with a data stream, you may require continuing assurance that no agent has tampered with the stream since the connection was opened.

It should be clear that efficient protocols for identifying agents and for signing documents and data are needed in Worldnet. Public key cryptosystems can provide these functions, but they are not yet widely available through standard network protocols (9). Cryptosystems can also meet the need to exchange private information.

In the past, most people associated authentication and privacy with military security and did not see justification for the cost of introducing these mechanisms into public networks. Recent incidents involving computer worms and viruses have changed this attitude. Most people now see that authentication and privacy are fundamental to coordinated action.

Authentication and coordination are not possible without a system of naming that associates character strings with agents and resources. Names are linguistic shorthand

for identities, and we associate our assessments of trust with them. Without names, Worldnet cannot function.

The naming system must be easy for people to use. For example, it should allow them to use short nicknames (aliases) for familiar agents. It should include directory services that provide the network names of agents or resources when given descriptions of their functions or characteristics. It should be hierarchical so that the authority to assign names can be delegated downward as far as possible. The telephone network uses a hierarchical system of "names" (telephone numbers) with country codes, area codes, prefixes, suffixes, and sometimes extensions within an organization; speed-dialing is a means for using short nicknames.

Network names must be independent of location. Otherwise a subsystem would fail the moment one of its constituent agents was moved to a new node in Worldnet. I distinguish between network names, used directly by agents, and location-dependent addresses: deep within network software are routines that map network names to binary addresses of nodes and routes. Without location-independent network names, Worldnet cannot be dependable.

The Research Internet uses a hierarchical system of location-independent names derived from organizational domains. For example, the Internet name "leiner@nsd.riacs.edu." identifies a particular user in the Networked Systems Division of the Research Institute for Advanced Computer Science in the education domain. The Worldnet naming system is likely to be a hybrid containing geographic elements and functional elements.

The naming system is useless without directories to assist agents in locating names given other information about the agents or resources sought. This idea is well developed in the telephone network, where standard, universal protocols quickly bring a user into contact with a directory-assistance operator who can provide a phone number. A similar concept is developing in the Research Internet, where domain name servers have been established to answer queries about agents within those domains.

Models for coordination of action, coupled with authentication, privacy, and naming, form the foundation for a dependable Worldnet. In normal courses of action, people will require additional functions. For example, a supplier will design, build, and distribute new systems. A consumer will locate suppliers and purchase their services. To provide support for these common actions of suppliers and consumers, at least five additional capabilities must be present in Worldnet: help services, aids for subsystem design, aids for subsystem assembly, spin-off to machines, and resistance to attack. These capabilities will depend on the domains to which they apply.

First, the ability to post and to gather information about available resources within Worldnet is an important characteristic of directory services. Agents can be designed to post notices of available resources and services in databases and directories. Other agents can gather information from these sources. Brokers and advertisers will use such functions heavily.

Second, the ability to design new subsystems of agents and resources must be supported by design aids that help represent the emerging plan and record the decisions made along the way. These aids include configuration management systems, version control systems, and manufacturing process systems, among others used in engineering.

Future design aids will add a capability for capturing declarations and other events that make up the "corporate memory" of large projects.

Third, the ability to assemble the components of a system and set them into motion requires tools we do not yet have. For example, we need a method of building a computation by specifying resources attached to Worldnet as parts that can be plugged together. New programming problems will arise from the massive numbers of components that will make up many computations.

Fourth, as routines performed by people become well understood, machines can be built that carry out those routines automatically. These new machines will be spin-offs of existing Worldnet functions.

Finally, the design of Worldnet must include security mechanisms that protect the network and its components from attack by malicious programs such as viruses and worms. Worldnet must have an immune system.

This discussion has suggested directions for research needed to realize Worldnet. The underlying theoretical basis must yield new distinctions, language, and notations for the domains of action in which people (and agents) will perform. The theory will be different in character from traditional scientific theory, which is quantitative and produces equations that can be used to make predictions in the world. The theory required for Worldnet is ontological -- it deals with the distinctions around which actions are possible, develops the language required for those distinctions and their relationships, and reveals the functional elements of architecture that will support people in their networks. Like traditional scientific theory, it must be rigorous, but many of its conclusions will not be directly testable by experiments. It will be inspired by many

disciplines, including computer science, mathematics, linguistics, social science, psychology, and behavioral science (5). It will focus not on what computers do, but on what people do with computers.

References

1. P. J. Denning. 1989. "Massive parallelism in the future of science." *American Scientist* 77, 1 (January). 16-18.
2. D. Langenberg (ed.). 1988. *Information Technology: The User's View*. National Research Council Report.
3. S. Berger, M. L. Dertouzos, R. K. Lester, R. M. Solow, and L. C. Thurow. 1989. "Toward a new industrial America." *Scientific American* 260, 6 (June). 39-47.
4. M. L. Dertouzos, R. K. Lester, R. M. Solow, and MIT Commission on Industrial Productivity. 1989. *Made in America: Regaining the Productive Edge*. MIT Press.
5. T. Winograd and F. Flores. 1987. *Understanding Computers and Cognition*. Addison-Wesley.
6. F. Flores, M. Graves, B. Hartfield, and T. Winograd. 1988. "Computer systems and the design of organizational interaction." *ACM Transactions on Office Information Systems*. April. 126-152.
7. A. W. Holt. 1988. "Diplans: A new language for the study of an implementation of coordination." *ACM Transactions on Office Information Systems*. April. 109-125.

8. A. W. Holt. 1988. "Coordination technology: Fulfilling CSCW's potential." Technical Report, Coordination Technology, Inc., 35 Corporate Drive, Trumbull, CT 06111.
9. P. J. Denning. 1987. "Security of data in networks." *American Scientist* 75, 1 (January-February). 16-18.

Steps toward Worldnet

Networking is widespread in the western industrial nations. In the United States, a large majority of scientists already have access to networks. Businesses are linking powerful workstations in local networks, and many use the telephone network to provide links with other organizations. Cellular telephone and fax are extending the network link to portable computers.

Several annual conferences that emphasize cooperative work have attracted large and increasing followings. These include Computer Supported Cooperative Work (CSCW), Computers and Human Interfaces (CHI), artificial intelligence, and graphics.

In 1987-88, NASA conducted a Telescience Testbed Pilot Program with 15 universities to explore the conduct of science as interactions with remote instruments supplying data streams to cooperating groups of investigators. Although telescience is currently focused on scientific uses of the space station, it can have general implications. Among the issues successfully explored were ways to use networking to augment scientific experiments. Follow-up studies are being conducted.

In 1987, NSF initiated the EXPRES (Experimental Research in Electronic Submission) project to experiment with preparing, communicating, and editing multimedia documents. This project, which was carried out by Carnegie-Mellon University and the University of Michigan, demonstrated the feasibility of electronic submission.

In 1987, the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET) prepared a report released by the president's science advisor calling for the US government to establish a high performance computing initiative that would include high-speed processors, high-bandwidth networking, software technology, and basic research and human resources. This report created a context for the Gore bill introduced in the Senate this year. FCCSET plans to release an implementation plan for the recommendations of the report.

In 1987, FCCSET spawned a subgroup called the Federal Research Internet Coordinating Committee (FRICC), consisting of representatives from the DoE, HHS, DARPA, NSF, and NASA. This group is developing a strategy to share the resources of the participating agencies and stimulate the creation of a commercial network capability beyond gigabit/sec transmission rates by 1996. Also in 1987, a

Coordinating Committee for Intercontinental Research Networking (CCIRN) was formed to coordinate network connections between North America (represented by FRICC and Canada) and Europe in support of the research community.

In 1988, the Open Software Foundation was formed as an alliance of eight major hardware and software manufacturers to develop and adopt common standards for operating systems, networking, windows, editing, and other elements that will be part of an international information infrastructure.

This year, a group of scientists met to develop recommendations to the NSF for a research agenda that would result in a national "collaboratory." The collaboratory would be a set of functions and practices, based on ubiquitous high-speed networks, that would enable scientists and laboratories to collaborate regardless of the distances separating them.

The High Performance Computing Act

In May 1989, Senator Albert Gore introduced a bill (S. 1067) that would establish a federal policy to maintain and increase America's leadership in high-performance computing, high-speed networking, software, basic research, and training of computer and computational scientists. The bill calls for the planning and implementation of a national high-performance computing program by action of the federal agencies, led by the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET). (FCCSET reports to the director of the Office of Science and Technology Policy, also known as the president's science advisor.) All funds for the program would be specified in the national plan and argued separately before Congress. The bill assigns roles to all the major government agencies involved with the current Research Internet. It calls for an annual progress report to the president from FCCSET.

Besides high-performance computing, the bill would establish four other initiatives, each with a lead agency and specific funding through 1994. The first is a National Research and Education Network. The network would link government, industry, and higher education; it would be phased out when commercial networks can meet the demand. It would have accounting mechanisms to charge individuals or groups, who would in turn be allowed to charge grants and contracts for network use. The second initiative is a national information infrastructure that would provide directories of users and resources, access to unclassified federal databases, rapid prototyping of computer chips by facilities connected to the network, access to other databases with assistance from artificial intelligence programs, and opportunities for international collaboration among researchers. The third initiative is the development of high performance software for a variety of scientific and engineering applications, with specific encouragement for approaches involving artificial intelligence. Software developers would no longer be required to turn over proprietary development systems on delivery of software. This initiative calls for the NSF supercomputer centers to continue to have the most advanced supercomputers developed by US manufacturers. The fourth initiative is continued promotion of basic research and education in computer technology, more training in computer and computational science, and encouragement of the development of technology transfer mechanisms.

1966 INTERMEDIATE EDITION



Mail Stop 230-5
NASA Ames Research Center
Moffett Field, CA 94035
(415) 694-6363

The Research Institute for Advanced Computer Science
is an Institute of the
Universities Space Research Association
The American City Building
Suite 212
Columbia, MD 21044
(301) 730-2656